


Capri: Achieving Predictable Performance in Cloud Spot Markets

Bogdan Ghit and Asser Tantawi

Mascots 2021

Virtual conference, November 5, 2021



Tech Lead | Sr. Software Engineer @ Databricks

- Data warehouse connectivity

Research Intern @ IBM T.J. Watson

- Spot market analysis

PhD @ TU Delft


- Scheduling in datacenters

<https://bogdanghit.github.io>

Cloud spot markets



On-demand instance



Spot instance
Added when bid > spot price

Announcing low-priority VMs on scale sets now in public preview

Posted on May 3, 2018

[Meagan McCrory](#), Senior Program Manager, Azure Compute



Google Cloud Platform Blog

Product updates, customer stories, and tips and tricks on Google Cloud Platform

Introducing Preemptible VMs, a new class of compute available at 70% off standard pricing

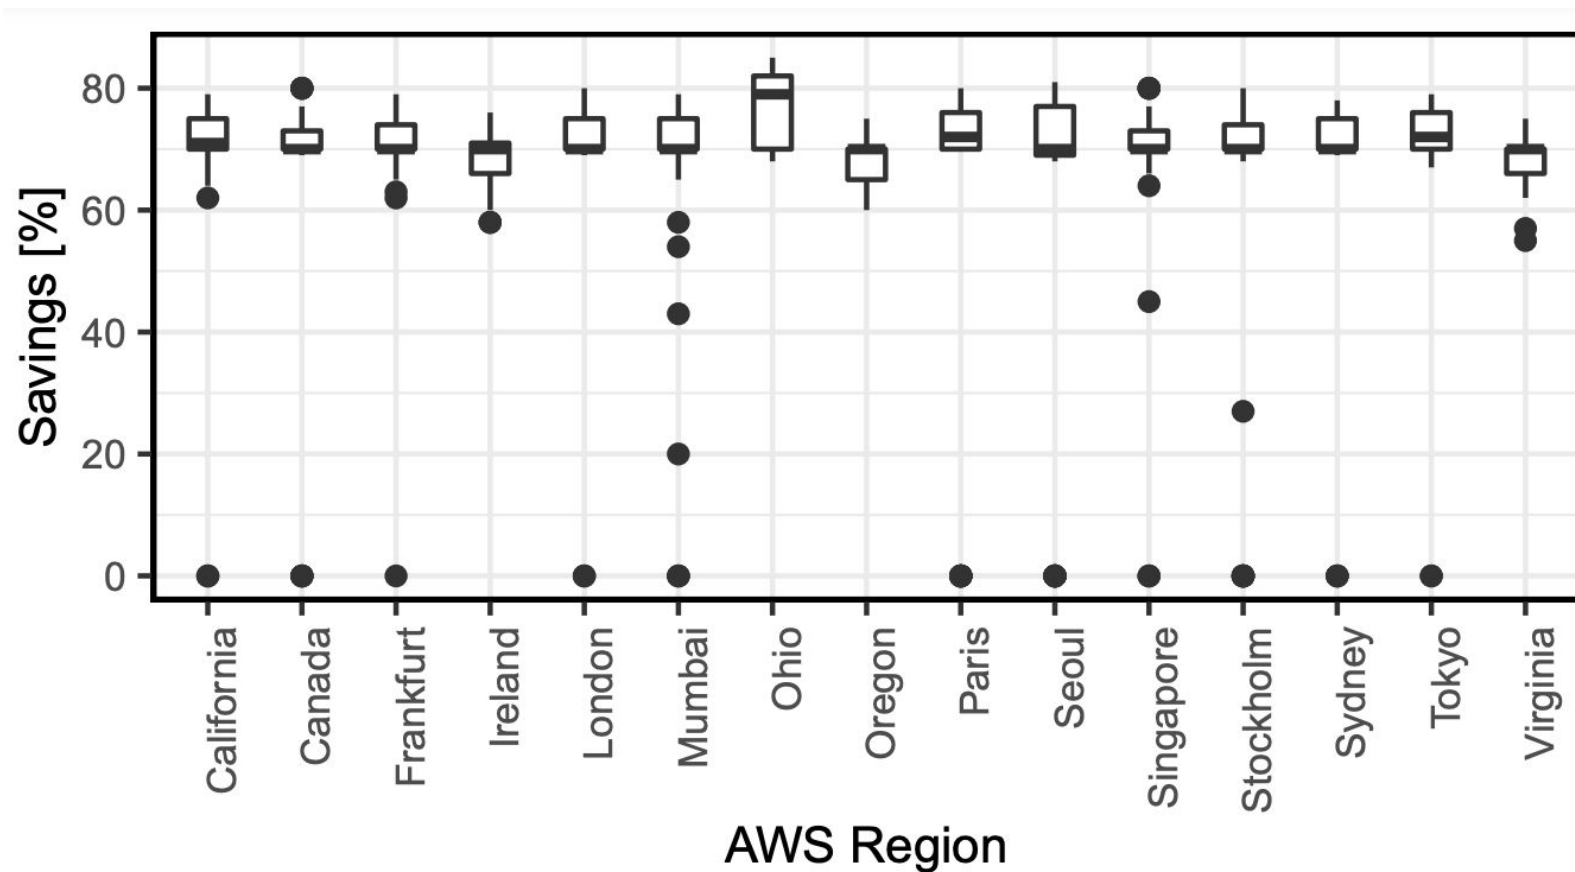
Monday, May 18, 2015

Announcing Amazon EC2 Spot Instances


Posted On: Dec 14, 2009

We are excited to announce the introduction of Amazon EC2 Spot Instances, a new way to purchase and consume Amazon EC2 Instances. Spot Instances allow customers to bid on unused Amazon EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price. The Spot Price changes periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. Spot Instances are complementary to On-Demand Instances and Reserved Instances, providing another option for obtaining compute capacity. If you have flexibility in when your applications can run, Spot Instances can significantly lower your Amazon EC2 costs. Additionally, Spot Instances can provide access to large amounts of additional capacity for applications with urgent needs. To learn more, please visit the [Amazon EC2 Spot Instances detail page](#).


Spot instances are attractive



... but they may disappoint users

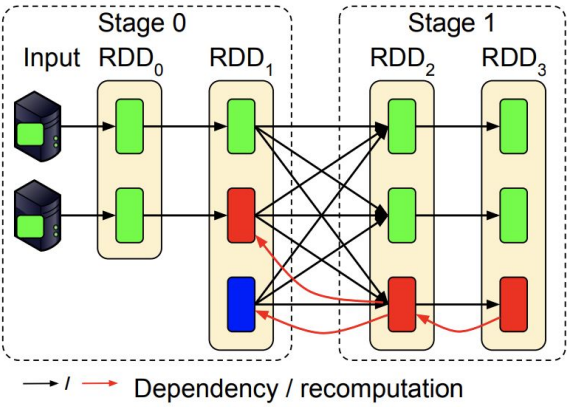
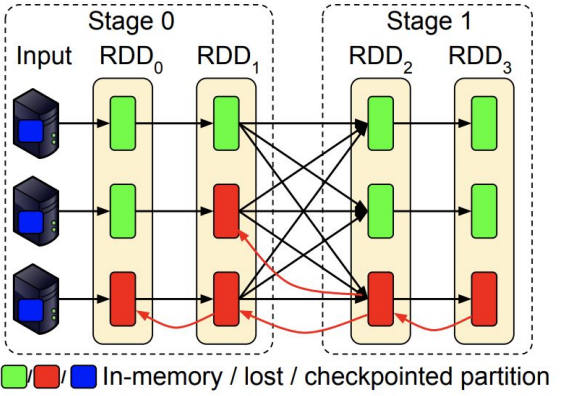


Cloud analytics applications are challenging



Flexible execution model → grow and shrink over time

Data dependencies → cascading recomputations



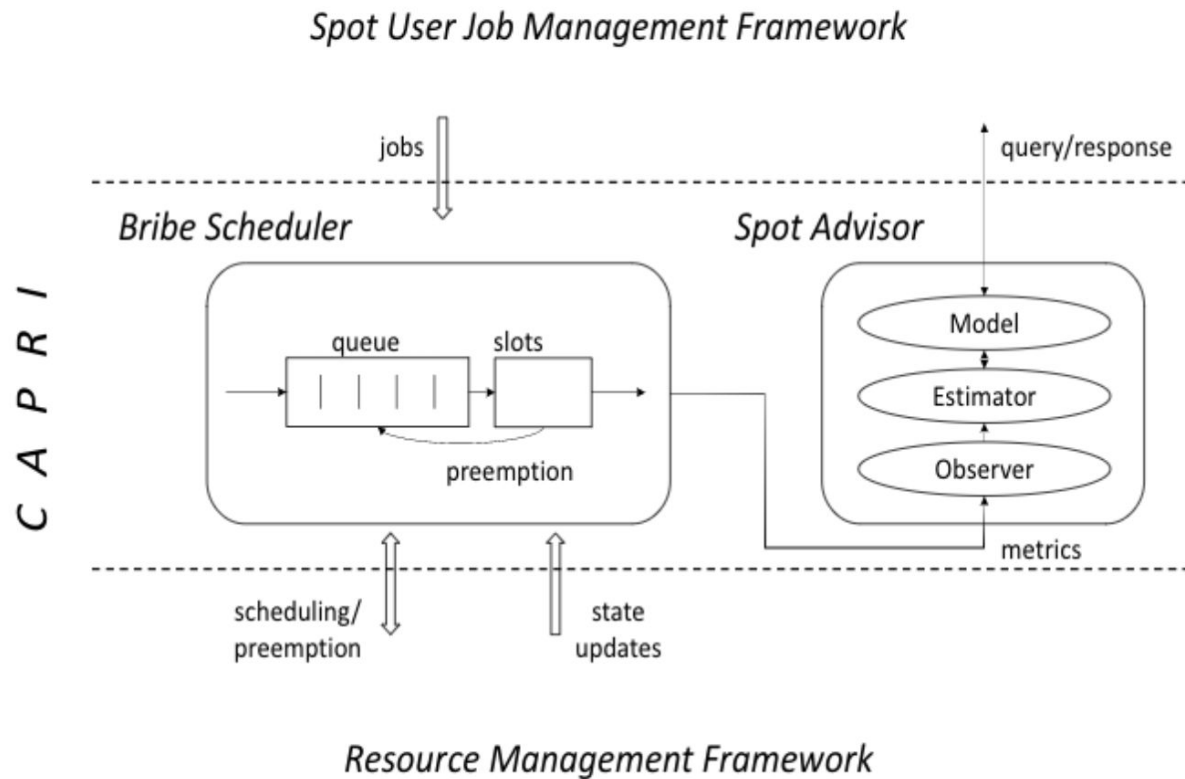
This work

Data analytics jobs suffer large slowdowns due to preemptions

Approach	Risk adversity	Costs	Performance
Bid the on-demand price	Low	High	Optimal
Periodic checkpointing	Medium	High	Suboptimal
Portfolio allocations	Medium	Medium	Variable

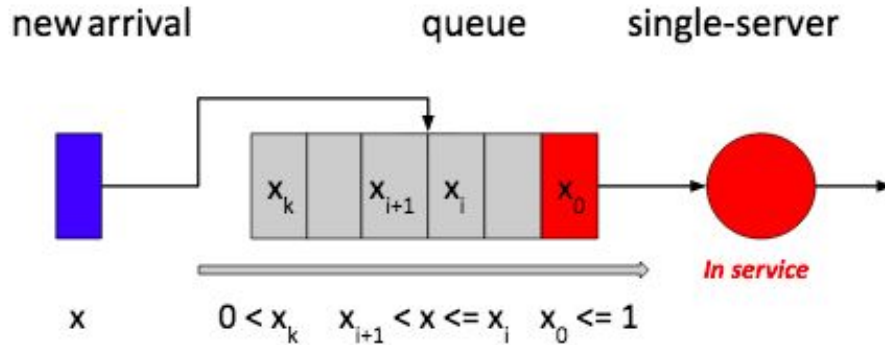
We want to set an expectation of the job slowdown as a function of the bid value

Overview of the Capri spot market



Macro-model: the bribing queue

M/M/1 bribing queue

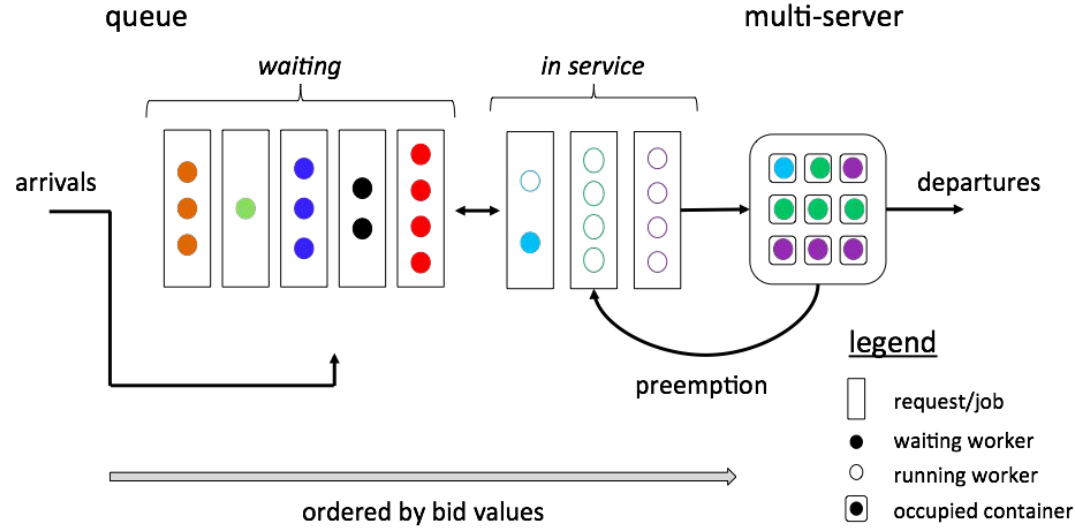


$$S(x) = \frac{1}{(1 - \rho(1 - B(x)))^2}$$

Kleinrock

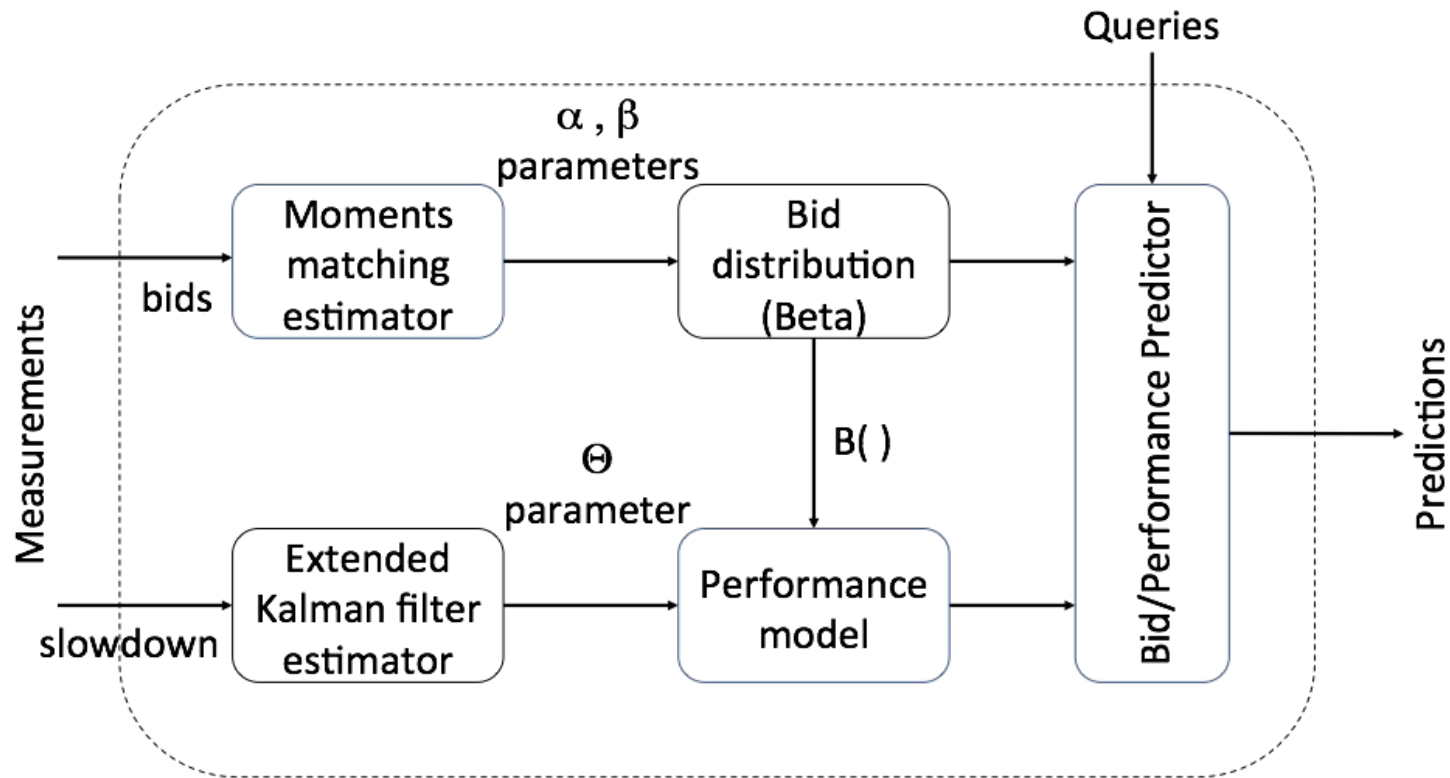
The bribe scheduler

Multi-worker, multi-server bribing queue

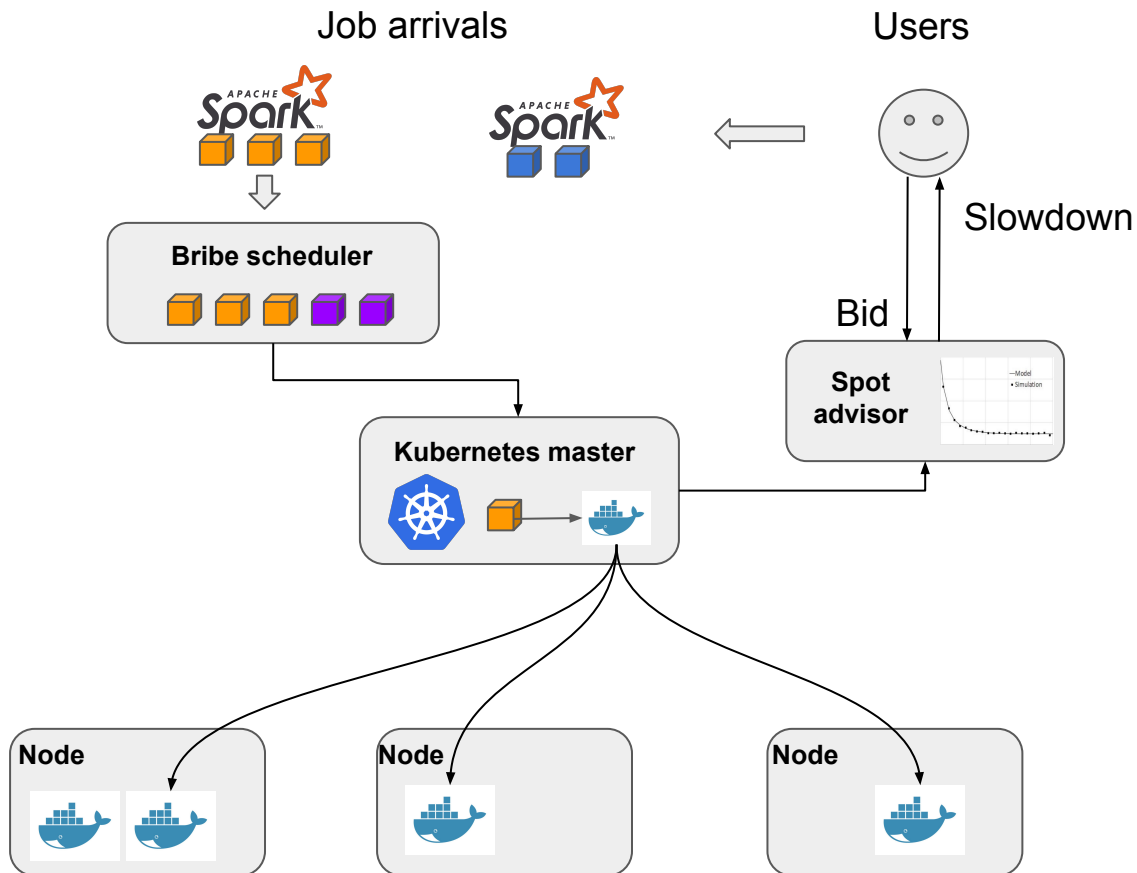


$$S(x; B, \Theta) = \frac{1}{[1 - \theta_0(1 - B(x))^{\theta_1}]^2}$$

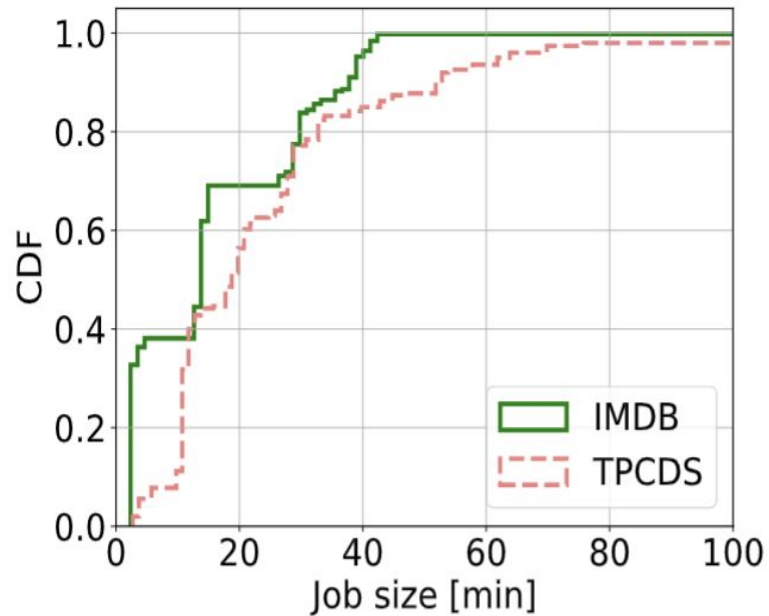
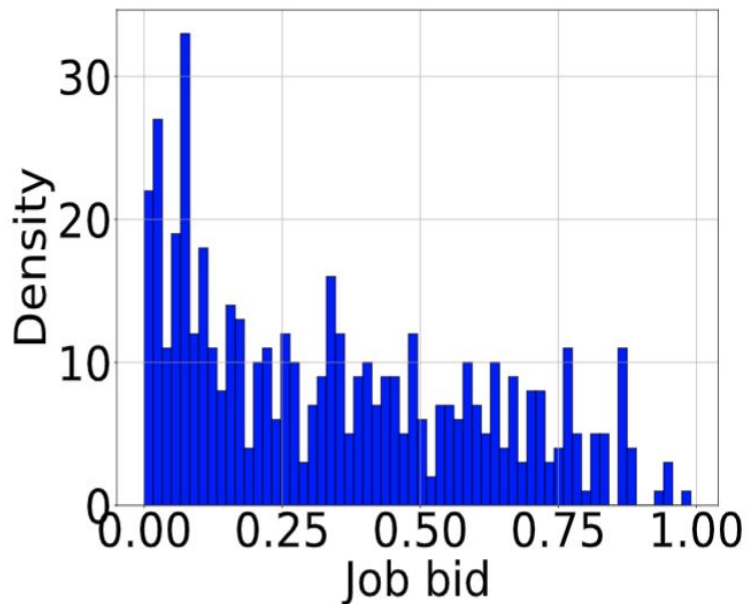
Prediction framework




System implementation




Experimental setup



Predicted job slowdown




IMDB

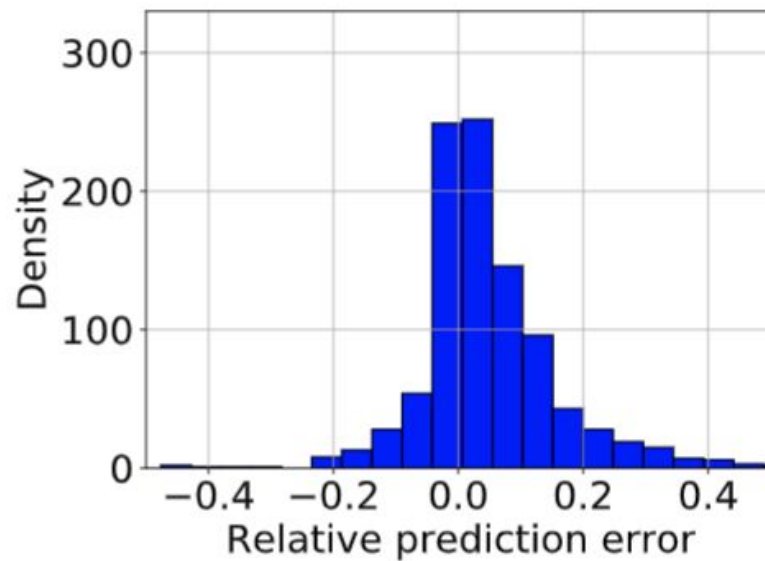


TPCDS

Prediction accuracy

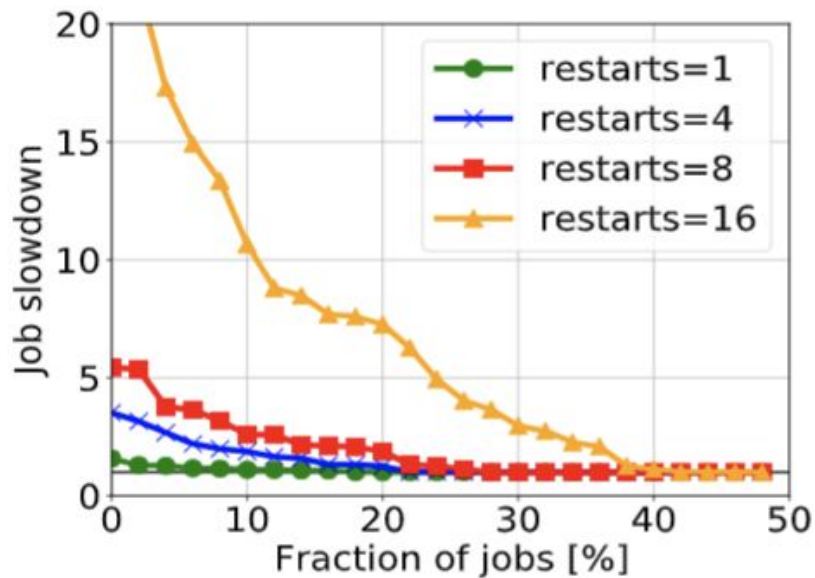


IMDB




TPCDS

Sensitivity analysis



IMDB



TPCDS

Conclusion

→ **Capri: alternative spot market with predictable performance:**

- ◆ Designed from first principles from a bribing queue macro-model
- ◆ Workload adaptation using Kalman filtering
- ◆ Extended with features to operate well in a real environment

→ **Summary of results:**

- ◆ Less than 3% median prediction error
- ◆ Strong correlation between observed and predicted job slowdown
- ◆ Relatively low sensitivity to the number of restarts